

# GPIC- GPU Power Iteration Cluster

Gustavo R.L Silva<sup>a,c</sup>, Rafael R. Medeiros<sup>b,c</sup>, Antônio P. Braga<sup>a</sup>, Douglas A.G. Vieira<sup>c</sup>

<sup>a</sup> *Electronic Engineering Department - UFMG, Av. Antônio Carlos, 6627, Belo Horizonte, Minas Gerais, Brazil*

<sup>b</sup> *Computer Engineering Department - CEFET-MG, Av. Amazonas, 7675, Nova Gameleira, Belo Horizonte, Minas Gerais, Brazil*

<sup>c</sup> *ENACOM, Professor José Vieira de Mendonça, 770, Belo Horizonte, Minas Gerais, Brazil*

---

## Abstract

This work presents a new clustering algorithm, the GPIC, a Graphics Processing Unit (GPU) accelerated algorithm for Power Iteration Clustering (PIC). Our algorithm is based on the original PIC proposal, adapted to take advantage of the GPU architecture, maintaining the algorithm original properties. The proposed method was compared against the serial and parallel Spark implementation, achieving a considerable speed-up in the test problems.

*Keywords:* GPU, Data Clustering, Power Iteration Cluster.

---

## 1. Introduction

The progressive incorporation of data collection and communication abilities into consumer electronics is gradually transforming our society, affecting labor relations, and creating new social habits. Consumer electronics, as well as industrial and commercial equipments are continuously incorporating functions that go far beyond their basic purpose [1].

The synergies between these physical and computational elements form the base of a profound transformation in the global economy. This new scenario expands current data acquisition and processing technology infrastructure, and significantly enhances connectivity, communication, computation, control technology and decision-making capabilities [1].

Many current real problems may involve terabytes/petabytes of data with hundreds/thousands of variables that tend to rise continuously if the current growth rate is maintained [2]. Exponential growth of data storage capacity in the worldwide network of devices is forecasted for the next years, therefore, new adaptive data-driven algorithms that consider the computational aspects of big data problems are demanded [3].

This may appear as a dilemma in the area, since it would be expected that a higher ability of sampling would improve representativeness and performance [4]. However, the reality is that, in spite of a higher representativeness, most current methods are not computationally capable to deal with such an increasing high dimension and volume [5].

The new big data scenario requires the investigation of new methods for data summarization [6], subspace representation [7], clustering [8], optimization, as well as pattern recognition and forecasting with lower computational complexity. The implementation of such methods requires scalable computing platforms capable to store and process massive and high dimensional data [9], [10].

Therefore, the need for real-time inference turns dedicated hardware solutions, such as GPUs (Graphics Processing Units), into an alternative to scale-up learning algorithms [11], [12]. In general, the main challenges for data set clustering in this context are the following: first, high dimensionality, due to the huge amount of sources generating and storing data; second, large volume of data, since data collected by the sources can also be huge due to local storage capacity.

This paper extends the Power Iteration Cluster (PIC) algorithm [13] to GPU, called here as GPIC. As presented in the results section, this implementation provides a considerable speed-up when compared with the serial and parallel PIC (Spark) implementation [14].

## 2. Cluster Analysis

A cluster is a collection of data objects which are similar to each other within the same group, or dissimilar to the objects in other groups [15]. The cluster analysis can be defined as: given a representation of a set of data points, find  $k$  groups based on a metric groups such that the similarities of the objects of the same group are high while the similarities of objects in different groups are low [15].

It aims at revealing the intrinsic structure of a given data set by forming statistically significant groups. Therefore, the aim is to separate a set of objects into groups where elements of the same group are similar to each other and the elements of different groups are has dissimilarities [16].

This analysis can be applied to: generate a compact summary of data for classification, pattern discovery, hypothesis generation and testing, outlier detection, dynamic trend detection, among others .

The Cluster Analysis presents some requirements and challenges [17]:

- Quality: ability to deal with different types of attributes such as numerical, categorical, text, multimedia, networks, and mixture of multiple types, discovery of clusters with arbitrary shape and ability to deal with noisy data;
- Scalability: clustering all the data instead of only on samples, high dimensionality and incremental or stream clustering and insensitivity to input order;
- Constraint-based clustering: user-given preferences or constraints like domain knowledge, user queries, interpretability and usability.

### 2.1. Graph-based methods

A graph-based methods first constructs a graph or hypergraph and then apply a clustering algorithm to a given partition [18]. A link-based clustering algorithm can also be considered as a graph-based one, since the links between

data points can be considered as links between the graph nodes [18]. Typical graph-based clustering algorithms are Chameleon [19], CACTUS [20], ROCK [21], Spectral Cluster [22].

#### 2.1.1. Spectral Cluster

Given a dataset  $S$  with  $n$  data points defined in  $R^m$ , it is defined an undirected and weighted graph  $G$ ,  $A$  is the affinity matrix of  $G$ ,  $D$  is the diagonal degree of the matrix  $d_{ii} = \sum_j^n a_{ij}$ , and  $W = D^{-1}A$  is a normalized affinity matrix.

To improve computational cost of Spectral Cluster algorithms, some authors have proposed new methods. Xiao et al.[23], proposed a new robust large-scale multi-view clustering method to integrate heterogeneous representations of large scale data.

Donghui Yan et al.[24], proposed a general framework for fast approximate spectral clustering in which a distortion-minimizing local transformation is first applied to the data. This framework is based on a theoretical analysis that provides a statistical characterization of the effect of local distortion on the mis-clustering rate [24].

Another approach is proposed by Frank and Willian [13] to find a very low-dimensional embedding of a dataset using truncated power iteration on a normalized pair-wise similarity data matrix. This embedding turns out to be an effective cluster indicator, consistently outperforming widely used spectral methods such as NCut on real datasets [13]. More details about PIC algorithm will be described on Methodology section.

### 3. Clustering data using GPUs

Recent advances in consumer computer hardware makes parallel computing capability widely available to most users. Applications that make effective use of the so-called Graphics Processing Units (GPU) have reported significant performance gains [25].

In CUDA (Compute Unified Device Architecture ) model, GPU is regarded as a co-processor which is capable of executing a great number of threads in parallel. A single source program includes host codes running on CPU (Central Processing Unit) and also kernel codes running on GPU. Compute-intensive and data-parallel tasks have to be implemented as kernel codes such a way to be executed on GPU [26].

In [27], the authors present an interesting study on how several data mining applications can be implemented using GPU in CUDA architecture [28]. In [29], a GPU version of k-means algorithm is presented. Another approach that uses k-means over GPUs is presented in [30], in this approach the authors applied k-means algorithm on a synthetic data set with a one billion data points and 2 dimensions. The GPU-version took about 26 minutes, and the serial version of the code took 6 days.

Jianqiang et al.[31], presented a GPU accelerated BIRCH (Balanced Iterative Reducing and Clustering using Hierarchies) version that can be up to 154 times faster than the CPU version with good scalability and high accuracy.

In [11] the authors present the G-DBSCAN, a GPU parallel version of the DBSCAN (Density-Based Spatial Clustering of Applications with Noise) clustering algorithm. This algorithm using GPUs, can be over 100 times faster than its sequential version using CPU.

## 4. Methodology

### 4.1. Power Iteration Clustering: An Overview

The PIC algorithm is a Spectral clustering technique. This algorithm proposes a method for computing the largest eigenvector of a matrix by the Power Iteration Method [32].

The Power Iteration Method [32] can be described like that, let  $W$  be a diagonalizable  $n \times n$  matrix with dominant eigenvalue  $\lambda_1$ . So, there is a nonzero vector  $x_0$  such that the sequence of vectors  $x_m$  [32]:

$$x_1 = Wx_0, x_2 = Wx_1, x_3 = Wx_2, \dots, x_k = Wx_{t-1} \quad (1)$$

tends to a dominant eigenvector of  $W$ .

$$|\lambda_1| > |\lambda_2| \geq |\lambda_3| \geq \dots \geq |\lambda_n| \quad (2)$$

Let  $v_1, v_2, \dots, v_n$  the corresponding eigenvectors. As  $v_1, v_2, \dots, v_n$  are linearly independent, they form a basis of  $\mathbb{R}^n$ . Consequently, we can write  $x_0$  as a linear combination of these eigenvectors:

$$x_t = W^t x_0, t \geq 1 \quad (3)$$

For

$$\begin{aligned} W^t x_0 &= c_1 \lambda_1^t v_1 + c_2 \lambda_2^t v_2 + \dots + c_n \lambda_n^t v_n \\ &= \lambda_1^t \left( c_1 v_1 + c_2 \left( \frac{\lambda_2}{\lambda_1} \right)^t v_2 + \dots + c_n \left( \frac{\lambda_n}{\lambda_1} \right)^t v_n \right) \end{aligned} \quad (4)$$

The fact that  $\lambda_1$  be the dominant eigenvalue means that each of the fractions, is less than 1 in absolute value. Like this,

$$\left( \frac{\lambda_2}{\lambda_1} \right)^t, \left( \frac{\lambda_3}{\lambda_1} \right)^t, \left( \frac{\lambda_n}{\lambda_1} \right)^t, \quad (5)$$

a sequence that tends to zero as  $t \rightarrow \infty$

Let  $\lambda_1 \neq 0$  e  $v_1 \neq 0$ ,  $x_k$  approximates a multiple of  $v_1$ , that is an eigenvector corresponding to  $\lambda_1$ , if  $c_1 \neq 0$ .

In the PIC [13] algorithm the Power Iteration Method [32] performs the update step:

$$v_{t+1} = \frac{W v_t}{\|W v_t\|_1} \quad (6)$$

A serial version of PIC [13] is presented in 1.

---

**Algorithm 1** Power Iteration Cluster [13]

---

**Input:**

$W$  Row-normalized affinity matrix  
 $k$  Number of clusters  
 $v_0$  Initial vector  
 $\epsilon$  Precision

**Output:**

$C$  Clusters  $C_1, C_2, \dots, C_k$

```
1:  $\delta_0 \leftarrow v_0$  ▷ Initialize the variation of eigenvector
2: for  $t = 0, 1, 2, \dots$  do
3:    $v_{t+1} \leftarrow \frac{Wv_t}{\|Wv_t\|_1}$  ▷ update the eigenvectors of  $W$ 
4:    $\delta_{t+1} \leftarrow |v_{t+1} - v_t|$  ▷ Update the variation of the previous and current
   eigenvector
5:   if  $|\delta_{t+1} - \delta_t| \leq \epsilon$  then ▷ Stop criteria
6:     break ▷ Terminate the estimation of eigenvectors
7:   end if
8: end for
9:  $C \leftarrow k\text{-means of } v_t$  ▷ Cluster eigenvalues
10: return  $C$ 
```

---

#### 4.2. Our Approach GPIC - Gpu Power Iteration Clustering

The authors of the PIC [13] algorithm have shared an implementation in MATLAB language [33]. With this code, it was conducted an experiment to evaluate the PIC behavior for a moderate amount of data. We used two synthetic datasets, two moons and three circles with  $n = 15,000; 30,000$  and  $45,000$  data points.

The results of the experiment indicate that the main bottleneck in the PIC [13] is computing pair-wise distance/similarity for all data points  $O(n^2)$  (in the worst case). On average consumes 88.61% of the PIC total time. The results of the experiment are presented in Table 1.

Table 1: Runtime (seconds) consume of PIC algorithm.( $m=2$ )

Dataset	n	A size	A Time [s]	PIC total time [s]	% total time
2 moons	15k	3.85GB	2.534	2.548	99.45%
3 circles			2.194,27	2.363,64	92.83%
2 moons	30k	16GB	18.666,39	21.680,05	86.09%
3 circles			18.064,53	21.091,17	85.64%
2 moons	45k	36.5GB	97.966,50	103.778,10	94.39%
3 circles			62.228,62	84.977,15	73.22%

To solve this bottleneck we propose a GPIC (GPU Power Iteration Clustering) approach. GPIC creates a set of CUDA kernels and can be called one after another to perform different steps of the PIC algorithm. In order to process large volumes of data, GPIC approach divided the data in chunks and these chunks are copied to the *device* iteratively.

GPIC Algorithm is described in 2 and the figure 2 presents the GPIC Algorithm execution flow. The main focus of GPIC is the Affinity Matrix and the Power Iteration step.

The first kernel **AffinityMatrix**( $S, p$ ) is launched to evaluated the Affinity Matrix  $A$ , where  $p$  is the number of threads. The  $A$  matrix is symmetric of size  $n \times n$ , as  $n$  is the size of the input. When launched this kernel has  $p$  threads and each thread is responsible for calculating a set of rows of the matrix  $A$ ,  $n_r$ . Thus, each element of the matrix  $A$  is indexed from  $id$  of each thread and current index  $i$ . This step has  $O(n^2/p)$  computational complexity, each CUDA core calculate parts of the  $A$  matrix, where  $p$  is the number of threads.

The second kernel **RowSum**( $A, p$ ) is launched in order to sum the lines of  $A$  matrix and the result of this operation is stored in the vector  $D$ . This step has  $O(n/p)$  computational complexity.

The third kernel **NormMatrix**( $A, D, p$ ) is launched to normalize the affinity matrix  $A$ , and the results is stored in  $W$  matrix. This step has  $O(n^2/p)$  computational complexity.



The fourth kernel **Reduction**( $D, p$ ) is launched to calculate the sum of the vector  $D$ . In the step 9 of GPIC algorithm 2, this operation is invoked again and the result is stored in  $\tau$ . This kernel has  $O(n/p + \log(n))$  computational complexity. Sequential reduction occurs in  $O(n)$  and in a thread block where  $N = P$  (processors) we have  $O(\log N)$  complexity. We used sequential addressing pattern suggests by Nvidia [34] because is conflict free. Figure 1 describe how this reduction kernel works.




Figure 1: Parallel Reduction: Interleaved Addressing [34].

The fifth kernel **Norm**( $v_{t+1}, \tau, p$ ) is launched to normalize the vector  $D$ , and results are stored in the  $v_t$ . In the GPIC algorithm step 102, this operation is launched again to calculate the new values of  $v_t$ . This step has  $O(n/p)$  computational complexity.

The last kernel **Multiply**( $v_t, W$ ) is launched to multiply the matrix  $W$  by the vector  $v_t$ . This step has  $O(n^2/p)$  computational complexity.

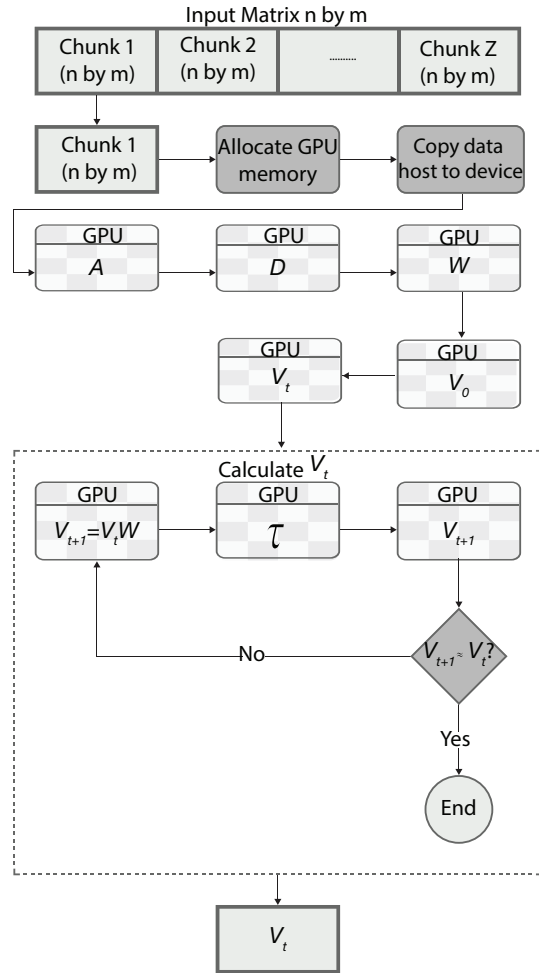


Figure 2: GPIC execution flow.

---

**Algorithm 2** GPU Power Iteration Cluster

---

**Input:**

$A$  Affinity Matrix  
 $k$  Numbers of clusters  
 $v_0$  Initial vector  
 $p$  Number of threads  
 $\epsilon$  Precision

**Output:**

$C$  Clusters  $C_1, C_2, \dots, C_k$

```
1:  $A \leftarrow \text{AffinityMatrix}(S, p)$  ▷ Affinity matrix  $A$  Kernel
2:  $D \leftarrow \text{RowSum}(A, p)$  ▷ Sum lines of  $A$  kernel
3:  $W \leftarrow \text{NormMat}(A, D, p)$  ▷ Normalize  $A$  kernel
4:  $v_0 \leftarrow \text{Reduction}(D, p)$  ▷ Sum  $D$  kernel
5:  $v_t \leftarrow \text{Norm}(D, v_0, p)$  ▷ Normalize  $D$  kernel
6:  $\delta_0 \leftarrow v_0$  ▷ Initialize the variation of eigenvector
7: for  $t = 0, 1, 2, \dots$  do
8:    $v_{t+1} \leftarrow \text{Multiply}(v_t, W, p)$  ▷ Multiply  $W$  by  $v_t$  Kernel
9:    $\tau \leftarrow \text{Reduction}(v_{t+1}, p)$  ▷ Sum  $v_{t+1}$  Kernel
10:   $v_{t+1} \leftarrow \text{Norm}(v_{t+1}, \tau, p)$  ▷ Update the eigenvectors of  $W$  Kernel
11:   $\delta_{t+1} \leftarrow |v_{t+1} - v_t|$  ▷ Update the variation of the previous and current
    eigenvector
12:  if  $|\delta_{t+1} - \delta_t| \leq \epsilon$  then ▷ Stop criteria
13:    break ▷ Terminate the estimation of eigenvectors
14:  end if
15: end for
16:  $C \leftarrow k\text{-means of } v_t$  ▷ Cluster eigenvalues
17: return  $C$ 
```

---

This GPU implemented PIC method converges to exact the same result of the original serial method, since the multi-thread explores the fact that the

operations are independent.

## 5. Results

### 5.1. Experiment I - Benchmark running time and speedup metrics

In this section the results of the proposed GPIC algorithm is compared to serial version of the PIC [13] and the parallel version of PIC implemented in Spark [14].

Experiments were conducted on a server containing two Intel Xeon E5-2620 (2 GHz, totally 24 cores) with 64 GB RAM and one k40m model NVIDIA card. This card k40m has 12 GB GDDR5 SDRAM and 2880 CUDA cores with 745 MHz.


Performance was evaluated using two synthetic datasets (two moons and three circles) with the matrices ranging in size from  $15.000 \times 15.000$  (3.85GB),  $30.000 \times 30.000$  (16GB) and  $45.000 \times 45.000$  (36GB), and each test case was run ten times. The average results are presented.

Table 2 shows the running time and speedup of PIC, PIC parallel (Spark implementation)[14] and GPIC, varying volume data size.


Table 2: Runtime (in seconds) and speedup comparison of PIC (serial version)[13], PIC parallel (Spark implementation) [14] and GPIC method on two synthetic datasets. The parameters for all experiments are maxiterations=3, precision=0.00001/n, m=2, and cosine similarity function.

Dataset	$n$	$A$ size	PIC serial Time [S]	Parallel PIC Time [S]	GPIC Time [S]	Speedup GPIC x PIC	Speedup Parallel PIC x PIC
2 moons	15k	3.85G	2.548	467	3,99	638,60	5,44
3 circles			2.363	320	4,03	586,51	7,38
2 moons	30k	16G	21.680	1.074	18,00	1.204,45	20,19
3 circles			21.091	1.063	18,03	1.169,78	19,84
2 moons	45k	36.5G	103.778	23.723	45,07	2.302,60	4,36
3 circles			84.977	22.976	45,90	1.851,35	3,69

This experiment results indicate a significant speedup gain for the GPIC approach which on average was 1,292.19 times faster than the original PIC in this two synthetic datasets. The parallel version of PIC implemented in Spark [14] get a speedup of only 10.15 times on average when compared with original serial PIC version. The figure 3 presents the run time with varying data sizes.



(a) Gpic  $\times$  PIC (Y axes in log scale)



(b) Parallel Pic on Spark  $\times$  PIC (Y axes in log scale)

Figure 3: Pic  $\times$  GPIC and Pic  $\times$  Parallel Pic on Spark

## 5.2. Experiment II - Subsampling to reduce input data matrix

One of the most important point observed in the previous experiment was that the increasing the amount of data can saturate both resources CPU and GPU memory.

To treat this problem, it was conducted another experiment to sample a portion of data and check the quality of the cluster with an external cluster validation indexes Adjusted Rand Index [35] and Jaccard Index [36].

This experiment considers four synthetic datasets (Cassini, Gaussian, Shapes and Smiley) with  $n$  equals 45.000 samples 5.

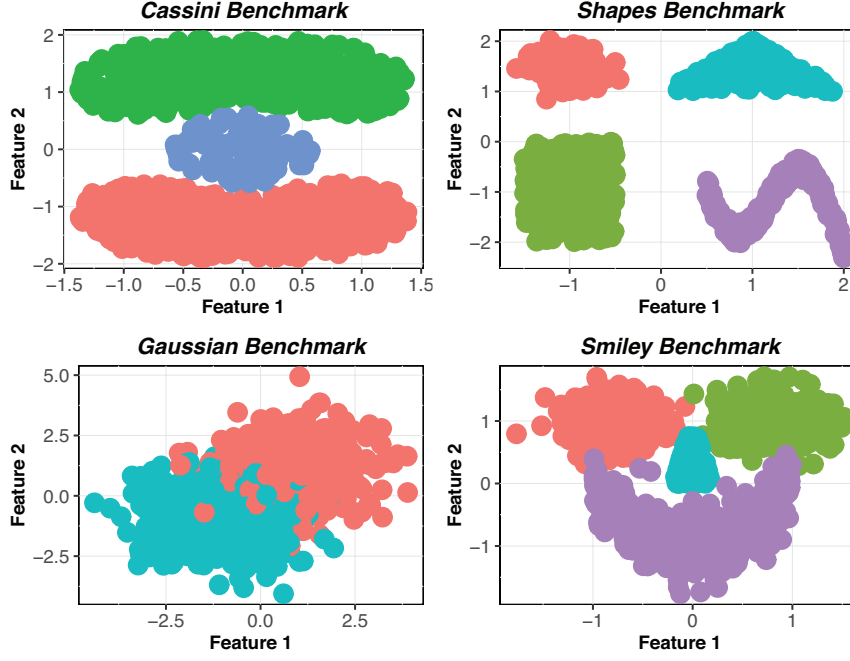


Figure 4: Four synthetic datasets used on experiment II.

Our approach to subsample the data considered the strategies bellow:

- size of data varies from 0.01% to 0.09% and 0.1%,to 0.9% accounting eighteen data samples;
- balanced classes;
- run ten times and collect the mean and standard deviation of Adjusted Rand Index [35] and Jaccard Index [36].

Through figure 5 it can be concluded that, for this four synthetic datasets the strategy to reduce the amount of data used in the  $X$  input matrix did not have a significant impact on the quality of GPIC when evaluated from the perspective of the Adjusted Rand Index [35] and Jaccard Index [36].

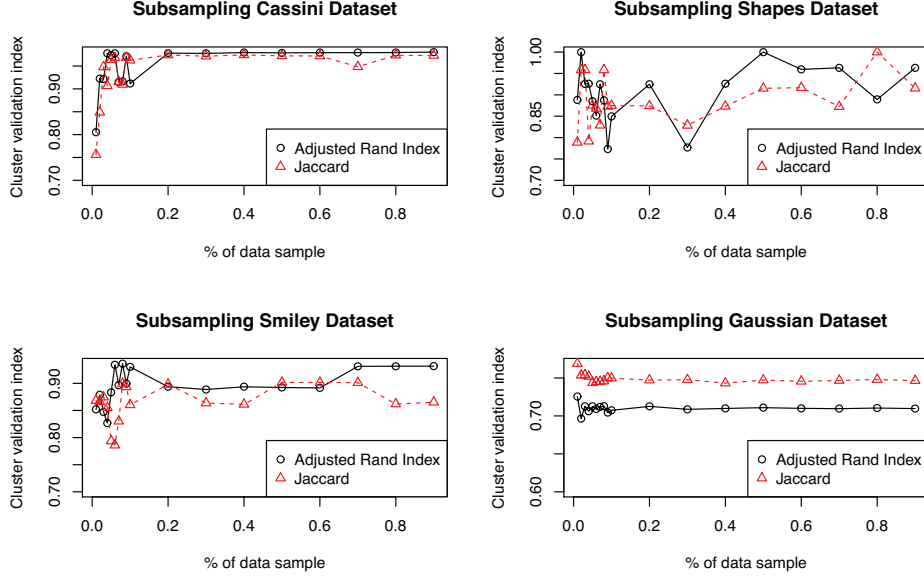


Figure 5: Four synthetic datasets used on experiment II.

## 6. Conclusion

This paper presents a new clustering algorithm, the GPIC, a GPU accelerated algorithm for Power Iteration Clustering. Our algorithm is based on the original PIC proposal, adapted to take advantage of the GPU architecture. The proposed method was compared against the serial and parallel Spark implementation, achieving a considerable speed-up in the test problems. Experimental results demonstrated that GPIC implementation has an excellent scalability.

It was analyzed the impact on cluster quality when the number of samples was reduced. With this experiment it can be seen that the cluster quality had no significant variations with data reduction. This experiment was very important to indicated a way when the data does not fit on CPU and GPU memory.

For future work we plan to use a version of GPIC algorithm on multiple GPUs boards to process big data problems. Multiple GPUs can compute faster (more GPUs equals faster time to a solution), larger (more GPUs means more



memory for larger problems) and cheaper (more GPUs per node translates to less overhead in money, power and space).

We think that on multi-GPU boards and sampling approaches to reduce the amount of data input, GPIC algorithm can achieves linear speedup, and significant performance improvement for a big datasets.

## References

## References

- [1] Z. Bi, L. Da Xu, C. Wang, Internet of things for enterprise systems of modern manufacturing, *Industrial Informatics, IEEE Transactions on* 10 (2) (2014) 1537–1546.
- [2] I. A. T. Hashem, I. Yaqoob, N. B. Anuar, S. Mokhtar, A. Gani, S. U. Khan, The rise of big data on cloud computing: review and open research issues, *Information Systems* 47 (2015) 98–115.
- [3] M. Chen, S. Mao, Y. Liu, Big data: A survey, *Mobile Networks and Applications* 19 (2) (2014) 171–209. doi:10.1007/s11036-013-0489-0. URL <http://dx.doi.org/10.1007/s11036-013-0489-0>
- [4] J. Fan, F. Han, H. Liu, Challenges of big data analysis, *National science review* 1 (2) (2014) 293–314.
- [5] A. Gandomi, M. Haider, Beyond the hype: Big data concepts, methods, and analytics, *International Journal of Information Management* 35 (2) (2015) 137–144.
- [6] J. Leskovec, A. Rajaraman, J. D. Ullman, *Mining of massive datasets*, Cambridge University Press, 2014.
- [7] L. Parsons, E. Haque, H. Liu, Subspace clustering for high dimensional data: a review, *ACM SIGKDD Explorations Newsletter* 6 (1) (2004) 90–105.

- [8] M. Bateni, A. Bhaskara, S. Lattanzi, V. Mirrokni, Distributed balanced clustering via mapping coresets, in: *Advances in Neural Information Processing Systems*, 2014, pp. 2591–2599.
- [9] J. Manyika, M. Chui, B. Brown, J. Bughin, R. Dobbs, C. Roxburgh, A. H. Byers, Big data: The next frontier for innovation, competition, and productivity. 2011, URL [http://www.mckinsey.com/Insights/MGI/Research/Technology\\_and\\_Innovation/Big\\_data\\_The\\_next\\_frontier\\_for\\_innovation](http://www.mckinsey.com/Insights/MGI/Research/Technology_and_Innovation/Big_data_The_next_frontier_for_innovation) 5 (33) (2014) 222.
- [10] A. Kaur, A. Datta, A novel algorithm for fast and scalable subspace clustering of high-dimensional data, *Journal of Big Data* 2 (1) (2015) 1–24.
- [11] G. Andrade, G. Ramos, D. Madeira, R. Sachetto, R. Ferreira, L. Rocha, Gdbscan: A gpu accelerated algorithm for density-based clustering, *Procedia Computer Science* 18 (2013) 369–378.
- [12] R. Bekkerman, M. Bilenko, J. Langford, *Scaling up machine learning: Parallel and distributed approaches*, Cambridge University Press, 2011.
- [13] F. Lin, W. W. Cohen, Power iteration clustering., in: *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, June 21–24, 2010, Haifa, Israel, 2010, pp. 655–662.
- [14] A. Spark, Clustering spark mllib, Apache Spark Site.  
URL <http://spark.apache.org/docs/latest/mllib-clustering.html#power-iteration-clustering>
- [15] A. K. Jain, M. N. Murty, P. J. Flynn, Data clustering: a review, *ACM computing surveys (CSUR)* 31 (3) (1999) 264–323.
- [16] O. R. Zaïane, A. Foss, C.-H. Lee, W. Wang, On data clustering analysis: Scalability, constraints, and validation, in: *Advances in knowledge discovery and data mining*, Springer, 2002, pp. 28–39.
- [17] A. K. Jain, Data clustering: 50 years beyond k-means, *Pattern recognition letters* 31 (8) (2010) 651–666.

- [18] G. Gan, C. Ma, J. Wu, Data clustering: theory, algorithms, and applications, Vol. 20, Siam, 2007.
- [19] G. Karypis, E.-H. Han, V. Kumar, Chameleon: Hierarchical clustering using dynamic modeling, *Computer* 32 (8) (1999) 68–75.
- [20] V. Ganti, J. Gehrke, R. Ramakrishnan, Cactusclustering categorical data using summaries, in: Proceedings of the fifth ACM SIGKDD international conference on Knowledge discovery and data mining, ACM, 1999, pp. 73–83.
- [21] S. Guha, R. Rastogi, K. Shim, Rock: A robust clustering algorithm for categorical attributes, in: Data Engineering, 1999. Proceedings., 15th International Conference on, IEEE, 1999, pp. 512–521.
- [22] A. Y. Ng, M. I. Jordan, Y. Weiss, et al., On spectral clustering: Analysis and an algorithm, *Advances in neural information processing systems* 2 (2002) 849–856.
- [23] X. Cai, F. Nie, H. Huang, Multi-view k-means clustering on big data, in: Proceedings of the Twenty-Third international joint conference on Artificial Intelligence, AAAI Press, 2013, pp. 2598–2604.
- [24] D. Yan, L. Huang, M. I. Jordan, Fast approximate spectral clustering, in: Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining, ACM, 2009, pp. 907–916.
- [25] J. D. Owens, M. Houston, D. Luebke, S. Green, J. E. Stone, J. C. Phillips, Gpu computing, *Proceedings of the IEEE* 96 (5) (2008) 879–899.
- [26] Y. Li, K. Zhao, X. Chu, J. Liu, Speeding up k-means algorithm by gpus, in: Computer and Information Technology (CIT), 2010 IEEE 10th International Conference on, IEEE, 2010, pp. 115–122.
- [27] J. Nickolls, I. Buck, M. Garland, K. Skadron, Scalable parallel programming with cuda, *Queue* 6 (2) (2008) 40–53.

doi:10.1145/1365490.1365500.

URL <http://doi.acm.org/10.1145/1365490.1365500>

- [28] M. Garland, S. Le Grand, J. Nickolls, J. Anderson, J. Hardwick, S. Morton, E. Phillips, Y. Zhang, V. Volkov, Parallel computing experiences with cuda, *IEEE Micro* 28 (4) (2008) 13–27. doi:10.1109/MM.2008.57.  
URL <http://dx.doi.org/10.1109/MM.2008.57>
- [29] A Parallel Implementation of K-Means Clustering on GPUs.  
URL <http://dblp.uni-trier.de/db/conf/pdpta/pdpta2008.html#FarivarRCC08>
- [30] R. Wu, B. Zhang, M. Hsu, Clustering billions of data points using gpus, in: *Proceedings of the combined workshops on UnConventional high performance computing workshop plus memory access workshop*, ACM, 2009, pp. 1–6.
- [31] J. Dong, F. Wang, B. Yuan, Accelerating birch for clustering large scale streaming data using cuda dynamic parallelism, in: *Intelligent Data Engineering and Automated Learning–IDEAL 2013*, Springer, 2013, pp. 409–416.
- [32] C. Lanczos, An iteration method for the solution of the eigenvalue problem of linear differential and integral operators, United States Governm. Press Office, 1950.
- [33] F. Lin, W. W. Cohen., Power iteration clustering, Nvidia Site.  
URL <http://www.cs.cmu.edu/~frank/>
- [34] M. Harris, Optimizing parallel reduction in cuda, Nvidia Site.  
URL [https://docs.nvidia.com/cuda/samples/6\\_Advanced/reduction/doc/reduction.pdf](https://docs.nvidia.com/cuda/samples/6_Advanced/reduction/doc/reduction.pdf)
- [35] L. Hubert, P. Arabie, Comparing partitions, *Journal of Classification* 2 (1) (1988) 193–218. doi:10.1007/BF01908075.  
URL <http://dx.doi.org/10.1007/BF01908075>

- [36] P. Jaccard, Nouvelles recherches sur la distribution florale, Bulletin de la Société Vaudense des Sciences Naturelles 44 (1908) 223–270.

This figure "GpicxPic.png" is available in "png" format from:

<http://arxiv.org/ps/1604.02700v1>

This figure "sparkXpic.png" is available in "png" format from:

<http://arxiv.org/ps/1604.02700v1>